AMENDMENTS

Amendments to the Specification

Please amend the title as follows:

A Method of Allocating Bandwidth in an Optical Network ~~Configurable Network Router~~

Please replace the paragraph on page 1, lines 5-13 with the following paragraph:

The present patent application is a continuation of U.S. Patent Application Serial No. 09/232,396, filed on January 15, 1999, Attorney Docket No. CIS0011US, entitled "METHOD OF ALLOCATING BANDWIDTH IN AN OPTICAL NETWORK" and ~~This application~~ is related to U.S. Patent Application Serial No. 09/232,395, filed on January 15, 1999, Attorney Docket No. CIS0009US ~~(Attorney Ref. No. M-7166 US)~~, entitled "A CONFIGURABLE NETWORK ROUTER~~A METHOD FOR ROUTING INFORMATION OVER A NETWORK~~," having A. Saleh, H. M. Zadikian, Z. Baghdasarian, and V. Parsi as inventors, and U.S. Patent Application Serial No. 09/232,397, filed on January 15, 1999, Attorney Docket No. CIS0008US ~~(Attorney Ref. No. M-7165 US)~~, entitled "A METHOD FOR ROUTING INFORMATION OVER A NETWORK~~A CONFIGURABLE NETWORK ROUTER~~, " having A. Saleh, H. M. Zadikian, J. C. Adler, Z. Baghdasarian, and V. Parsi as inventors~~, both of which are filed concurrently herewith~~. These related applications are assigned to ~~Monterey Networks, Inc.~~, Cisco Technology, Inc. the assignee of the present invention, and are hereby incorporated by reference, in their entirety and for all purposes.

Please replace the paragraph on page 8, lines 18-29 with the following paragraph:

According to ~~one~~ another embodiment of the present invention, a method is provided to allocate bandwidth from a first node to a second node in a optical network. The method begins by accepting a request from an end-user, who requests a virtual path between the first node and the second node. The first and second nodes are ones of a number of such nodes. Each one of the nodes is coupled to at least one other node by at least one of a number of optical links. The nodes and links form the optical network. The virtual path has a bandwidth requirement associated therewith. Next, the service provider determines an amount of bandwidth available between the first and the second nodes. The service provider then allocates at least a portion of the amount of bandwidth available between the first and second nodes equal to the bandwidth requirement, so long as the bandwidth requirement is not greater than the amount of bandwidth available between the first and second nodes.

Please replace the paragraph on page 9, lines 1 - 13 with the following paragraph:

According to yet another embodiment of the present invention, a method of allocating bandwidth in an optical network is provided. First, the service provider determines a bandwidth requirement of a requested virtual path between a first node and a second node. The first and second nodes are ones of a number of nodes. Each one of the nodes is coupled to at least one other node by at least one of a number of optical links. The nodes and links form the optical network. Next, a physical path between the first and the second nodes is selected from a number of such physical paths. The service provider then determines whether the physical path has enough available bandwidth to meet the bandwidth requirement of the requested virtual path. The steps of selecting a physical path and determining the available bandwidth for the physical path are repeated until either an acceptable physical path is found, or every one of the plurality of physical paths has been selected. If an acceptable physical path is found, the acceptable physical path is allocated.

Please replace the paragraph on page 11, lines 26-28 with the following paragraph:

1. Port Capacity growth: Router 100 includes, for example, a scaleable architecture which can provide, for example, (i) 250 or more nodes/network and (ii) at least 4096 ~~ports/nodes~~ ports/node, at a relatively low cost and high density.

Please replace the paragraph on page 12, lines 1-13 with the following paragraph:

2. Bandwidth management: The distributed management architecture of one embodiment of exemplary router 100 ~~according to the present invention~~ allows some or all nodes in the network to be managed from a single workstation. Provisioning a new connection is easily accomplished. Provisioning may be effected, for example, by selecting the source and destination nodes and specifying the required bandwidth and desired quality of service (QoS). An ~~optimized~~ QoS-based shortest-path first (SPF) ~~algorithm~~ path selection method is invoked to calculate the best route for the new connection. The ~~algorithm~~ QoS-based technique can take into consideration parameters such as existing

trunk allocations, network status, the priority and desired quality of the new connection, and other such criteria. This can be accomplished, for example, by sending one or more configuration requests to, and awaiting acknowledgment replies from, the nodes along the new ~~connections~~ connection's path.

Please replace the paragraph on page 12, lines 14-31 with the following paragraph:

3. Efficient and fast restoration: An exemplary network of two or more routers 100 as illustrated in Fig. ~~[NEW]~~ 1B preferably uses a mesh topology. Through the use of the routers 100, the network can be more efficient than existing ring topologies, especially when connections span multiple rings. This is possible because a single spare connection in a mesh network can provide protection for several different possible span cusps. In a ring-based network, however, spares can typically only protect against failures on their own ring. Furthermore, when connections span multiple rings, "dedicated" spare bandwidth must typically be allocated on every ring along the path. No other connections can share this spare bandwidth. Recent studies have found that mesh restoration typically requires only about 51% redundancy to yield 100% restorability, while a corresponding ring-based design typically require 330% redundancy for 100% restorability. One ~~emobodiment~~ embodiment of router 100 supports the restoration of a majority of network failures within less than 50 ms, thus eliminating an advantage that rings generally have over mesh topologies: fast restoration time. ~~An illustrated protocol~~ A protocol, such that according to the co-pending application entitled "A METHOD FOR ROUTING INFORMATION OVER A NETWORK" (as previously referenced), can be run on such a router and encompasses all aspects of the restoration process: alarm gathering, path implementation (including alternate path discovery), and path assurance. In cases where there is insufficient bandwidth to satisfy all failed connections, the protocol, in one embodiment, can use a quality of service (QoS) metric to prioritize the restoration sequence. In such embodiment, connections with the highest QoS are restored first, followed, in a descending order, by those with a lower QoS, until either all connections have been restored or all available bandwidth has been used.

Please replace the paragraph on page 13, lines 19-29 with the following paragraph:

In one embodiment, there are at least three types of processors in a router 100. The lowest level, level-3, resides on the line card, ~~or I/O module~~, also referred to herein as the I/O module and is responsible for all real time aspects of the processing of the physical protocol (e.g., SONET). In a SONET implementation, every level-3 processor is responsible for a single optical signal (e.g., an OC-48 signal) and, via a protocol processor, performs all required SONET/SDH section and line termination functions. The fast response time required from the level-3 processor makes a firmware implementation preferable. The firmware, which may be written in the "C" or "C++" programming languages, assembler, or other programming language, is preferably optimized for low latency and resource efficiency. Higher-level processing is implemented on a separate module, the shelf processor module, which is shared by several line cards.

Please replace the paragraph on page 14, lines 20-24 with the following paragraph:

~~A router~~ Router 100 provides yet another type of processor, referred to herein as a route processor. Such a processor is dedicated to the path/route discovery and restoration functions. The route processor is responsible for receiving failure indications from the line cards, calculating a new route for failed connections, and sending reconfiguration requests to all affected nodes, including its own.

Please replace the paragraph on page 14, lines 26-29 with the following paragraph:

In one embodiment, ~~a~~ router 100 is a multi-rack communications system capable of terminating at least 8192 signals and cross-connecting at least 4096 OC-48 signals. Such a router can be used, for example, as ~~a~~ SONET/SDH line terminating equipment (LTE) capable of terminating the Section and Line overheads of received OC-48 signals, and cross-connects those signals according to provisioned input-output mappings. Some of the terminated signals can optionally be protected using any of the common protection schemes (1+1, 1:1, and 1:N).

Please replace the paragraph on page 15, lines 4-8 with the following paragraph:

Overhead processing and generation is performed on the line card by a protocol processor ~~function~~. ~~This function~~ This protocol processor handles all aspects of the SONET protocol, including framing, insertion and extraction of embedded data channels, error checking, AIS detection, pointer processing, clock recovery, multiplexing/duplexing, and similar duties.

Please replace the paragraph on page 15, lines 10-21 with the following paragraph:

Fig. 2 is a block diagram of signal paths 200 within router 100. The primary signal paths in router 100 include one or more groups exemplified by groups 210(1)-(N), group matrices 212(1)-(N), and a main matrix 214. As depicted in Fig. 1A, groups 210(1)-(N), and group matrices 212(1)-(N) are shown ~~as~~ having receive and transmit sections. Groups 210(1)-(N) each include line cards 220(1,1)-(1,N), through line cards 220(N,1)-(N,N). Signals from line cards 220(1,1)-(N,N) are sent to the corresponding group matrix. In one embodiment, two sets of the group matrix cards, group matrices 212(1)-(N) and group matrices 216(1)-(N) are employed. Main matrix 214 is also mirrored in one embodiment by a redundant copy, a backup main matrix 218, which together form switching matrix 130. As shown in Fig. 2, the redundancy for group matrices 212(1)-(N) (i.e. group matrices 216(1)-(N)), is also provided on the transmit side.

Please replace the paragraph on page 16, lines 19-29 with the following paragraph:

In one embodiment, switching matrix 130 is an errorless, rearrangeably non-blocking switching network~~, and is discussed in greater detail subsequently~~. In one embodiment, switching matrix 130 is a 256x256 switching network that consists of three columns and 16 rows of 16x16 switching elements that allow any of their inputs to be connected to any of their outputs. Also, preferably a single copy of the matrix is housed in a single rack that contains three shelves, one for each column (or stage) of the matrix. Each shelf contains cards housing the 16 switching elements in each stage. The switching element itself may include, for example, a 16x16 crosspoint switch, with optical transceivers, and a microcontroller for controlling the crosspoint switch and providing operational feedback to the level-2 processor. Communications between the two processors may be carried, for example, over an Ethernet connection. The level-2 processor in turn communicates with the level-1 and route processors using, for example, a redundant Ethernet connection.

Please replace the paragraph on page 18, lines 19-29 with the following paragraph:

At the bottom of the hierarchy is what is referred to herein as a group matrix, or a Group Ethernet Repeater in a system using Ethernet communications, and depicted in Fig. 3 as group matrices 212(1)-(N) and 216(1)-(N). Each one of group matrices 212(1)-(N) and 216(1)-(N), also referred to herein as a hub, a repeater, or concentrator, is a physical layer device and preferably supports a star network topology, such as the IEEE 802.3 10BASE-T networking standard. The redundant connections from line cards 220(1,1)-(N,N) in each of groups 310(1)-(N) are connected to two repeaters that reside on two separate copies of the group matrix module. Preferably, each one of line cards 220(1,1)-(N,N) supports two network ports (e.g., 10BASE-T Ethernet ports). The two sets of four signals from each port pass through a relay that selects one of them for connection to the LAN for purposes of redundancy.

Please replace the paragraph on page 19, lines 6-19 with the following paragraph:

Fig. 3 also illustrates certain features of router 100 pertaining to the relationship between shelf switches 320(1)-(N) and 321(1)-(N), and groups 310(1)-(N). Groups 310(1)-(N) are again shown, with regard to the control functions thereof. In this depiction of groups 310(1)-(N), line cards 220(1,1)-(N,N) are shown as being attached to networking devices, indicated here as group matrices. Group matrices 212(1)-(N) and 216(1)-(N) may be, for example, multi-port Ethernet hubs running at 10Mbps. Each of line cards 220(1,1)-(N,N) feed signals into two of group matrices 212(1)-(N) and 216(1)-(N). For example, line card 220(1,1) feeds received information to group matrix 212(1) and group matrix 216(1) ~~which correspond to group matrices 212(1) and group matrices 216(1), respectively~~. Group matrices 212(1)-(N) and 216(1)-(N) each feed a signal into shelf switches 320(1)-(N) and 321(1)-(N) of Fig. 2. Shelf switches 320(1)-(N) and 321(1)-(N) are each controlled by a shelf processor (not shown) and communicate with one of the system switches (not shown, for the sake of ~~clarity~~ clarity).

Please replace the paragraph on page 19, lines 20-27 with the following paragraph:

Shelf switches 320(1)-(N) and 321(1)-(N) are the next higher level of the control hierarchy in router 100, and are located on the shelf processor module (exemplified by line racks (330(1)-(N)). Each copy of shelf switches 320(1)-(N) and 321(1)-(N) interconnects six connections from the three groups in each shelf, another connection from the shelf processor, and one connection from system switch ~~340(and 341)~~ 340 (and 341). Shelf switches 320(1)-(N) and 321(1)-(N) can be implemented, for example, using an 8-port Ethernet configured to handle 10Mbps Ethernet traffic and a single-port, dual-rate switch (e.g., 10Mbps/100Mbps Ethernet).

Please replace the paragraph on page 21, lines 1-8 with the following paragraph:

are defined with three of the groups occupying five slots each. In that embodiment, the other group, which occupies a single slot can be configured to house the shelf processor. Thus, in the embodiment shown in Fig. 4, I/O bay 400 contains line cards (LC) 410(1)-(N), group matrices ~~(SM)~~ GM 420(1)-(N), which are controlled by shelf processors (SP) 430(1)-(N), which are exemplary of line cards 220 (1,1)-(N-N), group matrices 212 (1)-(N) and 216 (1)(N), and ~~group matrices 320~~ shelf processors 320 (1)-(N) and 321(1)-(N), and shelf switches 440(1)-(N). It will be noted that the various line cards, group matrices, and shelf processors correspond to similar elements from previous figures.

Please replace the paragraph on page 21, lines 10-18 with the following paragraph:

A group is made up of line cards occupying a number of slots on a shelf. A slot is also referred to herein as a magazine In one implementation, the group is 20 line cards that occupy five slots. Four of the slots hold, for example, 16 ~~I/O modules~~ line cards at 4 per slot. The same slot can be used with a wide variety of I/O modules and in various configurations. One example of this flexibility, in a SONET configuration, is the ability to house an OC-192 I/O ~~modules~~ line card in the same space occupied by four OC-48 ~~modules~~ line cards. In fact, the slots in each group are not required to be of the same type or structure. This architecture provides ~~felxibility~~ flexibility to allow any combination of ~~modules~~ line cards to be installed in each slot.

Please replace the paragraph on page 21, lines 19-27 with the following paragraph:

The fifth slot in the aforementioned embodiment can be configured to accept line cards containing an optical switching matrix and a hub (e.g., an Ethernet hub). Preferably, two group matrix cards are employed, each containing a 2:1 optical reduction stage that "selects" working channels before the signals leave the shelf. In a 1+1 protection ~~sheme~~ scheme, the two inputs to the line cards are classified as active and protect channels. The working channel is one of the active and protect channels that is selected based on bit error rate or other criteria, and so implements a redundancy scheme. This prevents the standby ~~I/O modules~~ line cards from using any bandwidth on switching matrix 130.

Please replace the paragraph on page 22, lines 7-13 with the following paragraph:

Shelf processor module backplane signals include ~~the~~ reset signals, clock signals, hardware detect signals (e.g., card detect, copy present, and the like), ~~and~~ slot ID ~~signals and~~ signals, and slot communication signals (both low and high speed). I/O module line card backplane signals include reset signals, clock signals, communication signals, hardware detect signals, and slot ID signals. Group matrix module backplane signals include reset, clock signals, communication signals (both low and high speed), detection and hardware detect signals, and slot ID signals.

Please replace line 14 on page 22 with the following line:

~~(Line Card) I/O Module~~ Line Card

Please replace the paragraph on page 22, lines 16-22 with the following paragraph:

Fig. 5 illustrates the major components of one of line cards 220(1,1)-(N,N), exemplified in Fig. 5 by a line card 500. A line card, also referred to herein as an I/O module, integrates all the necessary hardware and software functions to properly terminate the physical layer. In a SONET implementation, ~~the I/O module~~ a line card terminates the transport overhead (Section + Line) of a full duplex OC-48 signal. Other components on this card provide a redundant optical connection to the switch matrix, and a communication channel to other modules in the system

Please replace the paragraph on page 22, lines 23-29 with the following paragraph:

Line card 500 receives optical signals from other network elements via a line-side optical receiver 505 and from the local router's system via a system-side optical receiver 506. Each of these receivers implements an optical-to-electrical (O/E) conversion function. Line card 500 transmits optical signals to other network elements using a line-side optical transmitter 510 and to the group matrices using a system-side optical transmitter 511. Each of these transmitters implements an electrical-to-optical ~~(O/E)~~ (E/O) conversion function. It will be noted that line-side refers to

Please replace the paragraph on page 23, lines 3-13 with the following paragraph:

Line-side optical receiver 505 is coupled to a protocol processor 520 which performs clock recovery multiplexing, demultiplexing, and SONET STE/LTE processing in both directions. Similarily, system-side optical receiver 506 is also coupled to protocol processor 520 to allow protocol processor 520 to receive optical signals. ~~In both transmit directions, the~~ The processed electrical signals from protocol processor 520 are coupled to the transmitters 510 and 511. The clock recovery functions are combined with demultiplexers and multiplexers to support reception and transmission of the optical data, respectively. The multiplexers serialize output data generated in protocol processor 520 by performing parallel-to-serial conversion on the parallel data. In contrast, de-multiplexers are used in protocol processor 520 to perform serial-to-parallel conversion on received data.

Please replace the paragraph on page 23, lines 14-18 with the following paragraph:

In order to add protection channels, line-side optical transmitter 510 is also coupled to a 1:2 broadcast unit 535. To receive such optical signals, optical receiver 506 is also coupled to a 2:1 selector 536 in order to select the working channel before the optical signals leave the shelf and thus prevent the standby ~~I/O modules~~ channel (also referred to herein as the protect channel) from using any bandwidth on switching matrix 130.

Please replace the paragraph on page 23, lines 19-24 with the following paragraph:

Protocol processor 520 is coupled to a bus 545. Protocol processor 520 interfaces the ~~I/O module~~ line card 500 to two copies of the matrix in a 1+1~~,~~ physical protocol. In a SONET implementation, protocol processor 520 provides both STE/LTE processing according to published industry standards ~~(e.g. Bellcore GR-253)~~. Also coupled to bus 545 are a memory 560 and a CPU 570. Memory 560 should be fast enough for efficient operation of CPU 570.

Please replace the paragraph on page 23, lines 25-29 with the following paragraph:

CPU 570 communicates with other of line cards 220(1,1)-(N,N) over a control bus (not shown) using a transceiver 580 that is coupled to CPU 570. Transceiver 580, is coupled to a transformer 585 which is coupled to a switch 590. Switch 590 is coupled to the control bus. Switch 590 implements a 1:1 protection scheme for transceiver 580 and couples CPU 570 to two independent ports on the backplane (not

Please replace the paragraph on page 24, lines 4-8 with the following paragraph:

Preferably, CPU 570 includes numerous integrated peripherals including embedded SCC channels (e.g. M-band communications) and an Ethernet controller (for example, to support communications with other system modules). In one embodiment, CPU 570 provides an onboard communications processor module (not shown) that handles time-critical aspects of the protocols supported.

Please replace line 9 on page 24 with the following line:

~~Group matrix Module~~ Group Matrix Module

Please replace the paragraph on page 24, lines 13-23 with the following paragraph:

Fig. 6. illustrates an exemplary group matrix 600, which is exemplary of group matrices 212(1)-(N) and group matrices 216(1)-(N). In the embodiment shown in Fig. 6, group matrix 600 includes a series of 2:1 path selectors (exemplified by selectors 610(1)-(N)~~)~~, broadcast units 620(1)-(N)), and a microcontroller 630 controlling these. Selectors 610(1)-(N) select one of two full-duplex optical signals and couple the selected signal to switching matrix 130. Selectors 610(1)-(N) and broadcast units 620(1)-(N) are grouped into pairs to form I/O channels 645(1)-(N). Microcontroller 630 communicates with other elements of router 100 via redundant transceivers (exemplified by transceivers 635 and 640). For example, microcontroller 630 can control selectors 610(1)-(N) and broadcast units 620(1)-(N) through commands received from the group processor.

Please replace line 24 on page 24 with the following line:

~~*Hub*~~ *Hub* **(no underline)**

Please replace the paragraph on page 24, lines 25-27 with the following paragraph:

One or more hubs are also provided to support communication between the group matrices and system switches in ~~a~~ router 100. In an Ethernet communications environment, the hub's functions are carried out primarily by repeater interface

Please replace the paragraph on page 25, lines 6-8 with the following paragraph:

The shelf processor module provides, among other elements, a shelf processor and switch that interconnect the LAN segments from the groups and the shelf processor to a port on the shelf switch (Ethernet switch ~~230~~ 730).

Please replace the paragraph on page 26, lines 1-7 with the following paragraph:

Shelf CPU 705 also accesses a memory 721 and a reset latch 722 over a CPU bus 725. Reset latch 722 supports reset of the ~~I/O~~ one or more line cards (not shown). Shelf CPU 705 is also coupled to an Ethernet switch 730. The network switch interconnects the lower speed inter-processor communication network segments in each shelf. In one embodiment, the network switch provides support for 10Mbps and 100Mbps segments. In one embodiment, an integrated bus master and slave interface~~,~~ allow multiple devices to be interconnected.

Please replace the paragraph on page 26, lines 16-21 with the following paragraph:

One embodiment of a system switch (or system Ethernet switch, in routers that communicate using Ethernet) capable of interconnecting at least 13 network segments in a switched configuration. In an Ethernet-based system, the system switch supports both 10 Mbps and 100 Mbps connections. The segments come from the shelf switching in the I/O shelf and the matrix switches, among others, and ~~allow~~ the system switch allows these elements to communicate.

Please replace the paragraph on page 26, lines 23-29 with the following paragraph:

A ~~switch fabric~~ switching matrix in router 100 is based on a rearrangeable non-blocking network. A ~~switch fabric~~ switching matrix, as described herein, consists of switch nodes arranged in a staged array. For a 256x256 ~~fabric~~ switching matrix, for example, ~~the network~~ switch matrix 130 consists of 48 nodes arranged in an array of 16 rows by 3 columns, with each column containing one stage of the switch matrix ~~130~~. All 48 nodes in the switch matrix ~~130~~ are substantially similar and consist of a ~~16x16 crossbar~~ 16x16 crossbar device that allows any of its 16 inputs to be connected to any of its 16 outputs, regardless of the current state of the crossbar.

Please replace the paragraph on page 27, lines 2-14 with the following paragraph:

Fig. 8 illustrates switching matrix 130 configured in the manner of the switch matrix just described. In one embodiment, switching matrix 130 employs a 256x256 matrix, an array of switching nodes 800(1,1)-(16,3), each of which is a 16x16 crossbar switch that allows any of the 16 input signals to be connected to any of its 16 outputs, regardless of the current state of the crossbar. In one environment, each of the interconnections between switching nodes 800(1,1)-(16,3) represent dual~~,~~ gigabit interconnections. As noted, the embodiment illustrated in Fig. 8 supports the switching of up to 256 inputs, shown as inputs 820(1)-(256). Inputs 820(1)-(256) are switched to one of outputs 830(1)-(256). Physically, each of the 48 switching nodes of this embodiment occupies a single slot in the matrix rack. The rack itself, which is shown in Fig. 9, is made up of three shelves (one per matrix column) that house the switch node cards (there are 16 such cards in every shelf) and six-shelf-processor cards (two per shelf).

Please replace the paragraph on page 28, lines 1-10 with the following paragraph:

The cross-connect information, i.e. input-to-output mapping, is written into the crosspoint switch by a local microcontroller which receives it from the local shelf processor over a high-speed connection. The three shelf processors in each rack receive such information from the node controller, which resides in a different rack. This hierarchy can be extended indefinitely. The crosspoint switch receives a high speed serial data from the optical receivers that perform optical-to-electrical conversion on the received optical signals. Data from the crosspoint switch is re-timed to synchronize the data ~~______~~ ~~**(delete blank line)**~~ with the system clock of router 100, using a clock and data recovery (CDR) unit, before being converted back into an optical signal that connects to the next stage of the matrix over fiber-optic cables.

Please replace the paragraph on page 28, lines 12-19 with the following paragraph:

Fig. 10 illustrates one of switching nodes 800(1,1)-(16,3) as a switching node 1000. Switching node 1000, in one embodiment, is a complete, strictly non-blocking, 16x16 OC-48 multi-stage crossbar matrix which allows any of its inputs to be connected to any of its outputs regardless of the current state of the matrix. A crosspoint switch 1005 is controlled by a local microcontroller (a microcontroller 1010) that also manages the optical transceivers, CDRs, and onboard SONET device. Switch node 1000 configuration is downloaded from microcontroller ~~1005over~~ 1005 over a low-speed bus.

Please replace the paragraph on page 29, lines 1-7 with the following paragraph:

receivers 1020(1)-(16) at switch input signals 1021(1)-(16). Crossbar switch 1010 also provides switch outputs 1022(1)-(16), which serve as the source of optical output signals for switch node 1000. Microcontroller 1010 is also responsible for detecting and reporting loss-of-signal (LOS) and out-of-lock (OOL) conditions from the optical receivers and CDRs, respectively. Microcontroller 1010 communicates with the shelf processor via transceivers 1060 and 1065 over a bus that carries asynchronous data over ~~the~~ a backplane (not shown).

Please replace the paragraph on page 35, lines 6-12 with the following paragraph:

The router 100 preferably runs one or more versions of the Unix operating system on the level-1 processor and the level-2 processors (in the I/O and matrix shelves). Level-2 processors preferably run a real-time version of the Unix operating system (OS). Other processors (e.g., level-3, route, ~~quad~~ group, and matrix-node processors) preferably run a single task that does not require the services of an operating system or kernel. While Unix operating systems are described herein as being preferable, any one ~~or~~ of a number of operating systems may be used.

Please replace the paragraph on page 35, lines 14-27 with the following paragraph:

The system controller is responsible for overall system management and control. The system controller uses a variety of protocols to communicate with other nodes in the network, including the operating system (OS). Some of the protocols satisfy specific requirements (e.g. in a SONET based system, the transfer of OAM&P message across the SONET/SDH communications channels DCC), while others implement features, or functions, that are not part of the physical protocol used. To facilitate these functions, every router ~~(one router, two, etc.)~~ in a network is assigned an ID that uniquely identifies it within the network. The ID can also serve as a priority metric that determines the node's level within the hierarchy. However, the network can be configured to allow the user to override this by manually assigning priorities to network nodes. The system controller supports a number of tasks that perform management, control, and routing functions, including resource management, OS interfacing, various network protocol servers, and operations, control, and intermediate system services.

Please replace line 17 on page 36 with the following line:

~~I/O Module (Line Card)Processor~~ *(Line Card)Processor*

Please replace the paragraph on page 36, lines 18-26 with the following paragraph:

The ~~I/O Module~~ line card terminates an input signal from one of the other nodes in the network. For example, in a SONET-based implementation, a single SONET/SDH OC-48 signal is terminated by ~~an I/O module~~ a line card, although other signal levels (OC-192, OC-12, and so on) may be supported. In one embodiment, the software consists of two threads, one that runs in the background and is responsible for non-time critical tasks. The other thread, which runs at the interrupt level, is responsible for all real-time aspects of the software, including limited overhead processing, alarm detection and forwarding, and fault detection and recovery. The ~~I/O module~~ line card processor maintains a copy of its firmware and startup code onboard.